CHAPTER

10

DISASTER PLANNING

After reading this chapter and completing the exercises, you will be able to:

- Describe the primary methods and devices used to back up critical data
- Identify methods for creating a high-availability network
- Explain the contents of a spare-parts kit for efficient hardware repair
- Identify key areas for SNMP monitoring
- Determine key server management and disaster recovery strategies for preserving system uptime

No matter how diligent you are at purchasing the best hardware possible and configuring it correctly, network failures are going to occur, either on a small scale when network equipment fails or on a large scale as in a natural disaster. Returning users to a state of productivity as quickly as possible after one of these failures is, of course, the goal. Being able to do so, however, requires a great deal of planning. If you rely upon sudden genius in the hour of need, you are likely to be disappointed.

One of the most important elements of a disaster recovery plan is a tape backup system. The administrator must not only know how to implement various types of backup strategies, but also physically store the tapes in the most appropriate locations. Besides the risk of losing data, the administrator must ensure that network resources are available, not only for normal productivity but also in case a server fails. This is where redundancy solutions such as clustering are critical for servers. Hopefully, you will be able to monitor the servers and predict when an impending failure is near. Recall from Chapter 9 that SNMP can alert you of such problems. Be sure to set SNMP thresholds so that you are notified well in advance of potential problems. Should problems occur that require only hardware replacement, be sure to keep compatible replacement hardware on hand in a spare-parts cabinet. In Chapter 2, you learned about the hazards that dust and heat pose to the server, and you should be certain to continue regular maintenance to guard against such problems.

In the worst-case scenario, a disaster occurs that physically obliterates the site. In this case, you must have a disaster recovery plan that reestablishes operations as quickly as possible, and probably involves one or more alternate sites.

BACKUPS

Although fault-tolerant hardware and devices can reduce total system downtime, the greatest potential loss in the event of a system failure of any kind is the loss of data. Even during normal day-to-day operations, data can become corrupt or lost. For example, users sometimes accidentally overwrite their own files with blank data or perhaps an application crashes while a file is open, resulting in a corrupt file. Having that data available for restoration is the highest-priority item in the preparation for disaster recovery.

The basic theory behind backing up data is simple: Never put yourself in a position where critical data for the survival of your network or business is permanently unavailable.

There are many backup devices from which to choose, and several backup methods that can be employed, but the goal is to never be without your mission-critical data. You will also have to decide who is going to be responsible for accomplishing the backups, and to what level these backups need to be secured. Which devices and methods you choose are not as important as the choice to develop and deploy a backup strategy.

Hewlett-Packard cites industry experts who estimate that the quantity of data increases 50% annually. If these figures hold true for your organization, then the need for accurate backup and restore operations is obvious, as is the need to select the appropriate backup hardware and strategies.

Types of Backup Hardware

Factors to consider when you purchase a backup unit include:

- The amount of data you need to back up
- Whether your backup software supports the unit
- The amount of money you want to spend
- The amount of time it takes to back up and restore data

You also need to consider the value of your data: Put a price tag on your data and on the time you'd need to rebuild the data if you lost it. Be willing to spend as much money on a device as some ratio of what you calculate your data to be worth.

There are many different types of backup media and devices, from removable media cartridges to CD-ROMs and magnetic tape. You need to determine the best balance

between speed of backing up the data, speed of restoration, the quality and function of supported backup programs, and storage media requirements. Then, choose a unit that provides the technology you need to back up and restore data efficiently and effectively within the parameters of your situation.

For example, some products use optical disks, which is a great solution for **Hierarchical Storage Management (HSM)** in which infrequently used data is moved from fast, expensive hardware (hard disks) to slower, less expensive media such as optical disks or magnetic tape. (An **optical disk** is any disk written and read by laser, including CD-R, CD-RW, DVD, and so forth.) However, optical disks are a poor solution if you want to back up large amounts of data quickly. In that case, you will usually turn to the ubiquitous and cost-effective tape drive. Tape drives offer the lowest cost per megabyte of storage, the greatest degree of flexibility compared to comparable-capacity devices, and are very portable.

The choices available in backup devices have greatly expanded over the last few years due to advances in technology. Table 10-1 introduces the most common backup solutions in the server realm: tape and CD laser devices.

Table 10-1 Common Server Backup Devices and Media Types

Device (Category)	Media Type	Pros	Cons
Tape drives (magnetic)	DAT, DDS, DLT	Inexpensive media, faster transfer	Expensive drives
CD (laser)	CD-ROM, CD-R, CD-RW	Inexpensive media	Limited software support for drive interface, not as flexible in terms of recording data as other solutions, slower than tape

Tape devices are generally the most popular server backup devices. Some devices support multiple tape formats, and others support only one. Some are very expensive, and some are not. The faster units are more expensive, as are units that have larger media capacities and/or automated handling of multiple media. One factor that often forces you to choose the large, more expensive drives is the backup window. The **backup window** is the optimal period of time in which you can perform a backup, and is usually when most files are closed. If your organization has a very small backup window, you need a very fast and expensive tape backup solution.

Tape devices are known for their long-lasting performance between hardware failures, which is partly attributable to the reliable tape drive mechanics and robotics that some systems include. Despite this reliability, you should research a vendor's service contract to ensure that the contract includes same-day on-site or overnight cross-shipping service if the unit fails.

Competition in the tape backup arena has created an abundance of dubious performance claims. You need to determine the type of files the vendor used to test the drive and the type of test the vendor ran, because many claims are a result of using totally compressible files. These tests don't take into account that many file types can't be compressed—so the tests don't simulate real-world scenarios.

Whichever type of tape drive you install, consider the importance of the tape drive driver. Plug and Play operating systems such as Windows 2000 and Linux might automatically detect hardware and install drivers of their own for the tape device. However, there might be features of the device that do not become available until you install the driver directly from the manufacturer. For example, many tape drives include hardware compression in which tape drive electronics compress data as it writes to tape, relieving the compression burden from the processor. However, unless you use the vendor drivers, the feature might not be available and you would then have to fall back to software compression utilizing the processor.

Be sure to upgrade tape drivers when upgrading from one drive to another. For example, upgrading from a stand-alone tape device to an autoloader without also upgrading the driver will likely result in a unit that will not back up at all or might only back up to a single tape.

No matter what type of drive you choose, versions of most backup products are available with a choice of SCSI, EIDE, or parallel port interfaces, except for the very high-end products, which all have SCSI interfaces. Additionally, some older, low-end products, particularly tape drives, also come in versions that connect to the floppy disk controller. Parallel port and floppy disk interfaces will not be considered further in this discussion of server backup solutions.

Departmental servers that might not require extremely high-capacity backup solutions might only require a single stand-alone EIDE tape device. These devices provide reasonably good performance at up to about 20 GB of compressed data. For the best backup performance, especially for use on higher-end workstations and certainly for all servers, the interface option of choice is SCSI. When assessing the cost of your backup solution, you'll need to add the price of a good-quality SCSI interface card for your server if you don't have one already, though most high-end workstation PCs and servers are already based on SCSI disk subsystems and therefore won't require an additional interface. In either case, check the type of physical SCSI connector fitted to a drive before buying to make sure you won't need additional cables or adapters. An additional benefit of SCSI is system responsiveness. With EIDE tape drives, some tape operations, such as cataloging the contents of the tape or re-tensioning, will temporarily dominate the system and affect responsiveness. Because of the nature of SCSI, you can perform nearly any tape function with minimal impact on overall system responsiveness.

If it is not cost prohibitive, take performance a step further and choose a tape solution that is Fibre Channel compatible.

Automated Tape Solutions

For a large-scale enterprise, none of the backup solutions outlined so far would be sufficient. Instead, you should probably investigate a backup technology that can automatically change tapes in a largely unattended fashion. Often large and extremely expensive, these automated tape devices are typically tape libraries that use autoloaders (see Figure 10-1). A **tape library** is a self-contained tape backup solution that is preloaded with several tapes. Most tape libraries include **autoloaders** to automatically load and swap tapes. In addition to automatic tape rotation, these devices can automatically clean tape heads.

Figure 10-1 A tape library using an autoloader

Automated backup solutions also mitigate potential human error. What if someone backed up Server_A and accidentally marked the tape Server_B? Restoring Server_A data might never occur unless someone manually goes through all the tapes searching for the data. Automated backup solutions won't write a label for you, but once you place the first set of tapes in the drives and properly configure the software, you can remove the human error factor from the tape backup. Plus, backups can be a tedious chore fraught with misplaced tapes, unmarked or mismarked tapes, incompatible tapes, and so on.

Many of the problems listed above relate to multiple persons performing backups. Make sure that only designated individuals perform the regular backups and that a written, logical procedure is defined.

Also, a properly configured tape library takes only a fraction of the time otherwise required with a manual backup. This becomes more of an issue as more data centers operate around the clock instead of during business hours, because the available backup window grows smaller.

Another big advantage of automated backup libraries is the extensive **online retention period (OLRP)**, the period for which data can be restored from tape without manual

intervention. Because a library can store several generations of backups at once, if you want to restore a file from, say, five weeks ago, you don't have to rummage through the storage cabinet to find the right tape. Instead, you can use the tape library software to locate and restore from the correct tape for you.

You might not have an automated backup device available, but at the very least, use a single tape drive to schedule nightly backups. Once you configure the scheduling software, all you have to do is replace tapes as necessary. I know somebody at a fast-growing organization who apparently got frustrated with the daily manual backup chore and stopped making tapes altogether. This became a big problem when a catastrophic server crash necessitated recovery using a recent tape backup—*which didn't exist*.

At the highest level, you'll see tape libraries that can support dozens of drives and hold several hundred cartridges (see Figure 10-2). At that level, you will probably not configure the tape solution yourself; you'll tell vendor representatives what your objectives are and they will configure the library for you or provide guidance.

Figure 10-2 The StorageTek L700 can store up to 678 tapes and has up to 20 drives

Types of Backup Tape Media

Tape media fall under a number of standards. The most common for smaller tape devices is probably the **Quarter Inch Cartridge (QIC)**, which as the name implies is a quarter inch in width. QIC cartridges have evolved over the years, starting at about 20 MB and progressing to 80 MB. The QIC Wide (8 mm) tape was developed to squeeze more data into tape cartridges, but you probably won't see QIC or QIC Wide in use on servers anymore. However, you should keep some drives on hand that can at least read this format, as there are about 200 million QIC cartridges worldwide and it might be necessary to restore an older archive. In 1994, a 3M spinoff (Imation) created a new QIC standard known as **Travan**. A Travan drive can reach compressed capacity of about 20 GB and is useful for home, small office/home office (SOHO), or small departmental backups. Because Travan drives accept the QIC format, they are usually backward read compatible with preceding QIC standards.

Travan NS (network series) (see Figure 10-3) is the most recent implementation of the Travan standard and addresses two main issues: compression and verification.

Figure 10-3 The Travan NS tape cartridge

Before Travan NS, data compression was always software based, taxing the processor and severely limiting other server functions. I recall when using early QIC and Travan technology that I often could not perform the simplest tasks while a backup was taking place. If I tried to do too much, the system would usually hang. Travan NS tapes and drives offer hardware compression instead, relieving the compression burden from the CPU.

Any tape backup job still affects system responsiveness regardless of the tape format and drive technology. Whenever possible, perform backups during off-peak times. This helps to assure that fewer files are open, which might prevent them from being backed up and cause inconsistency when comparing them against a backup verification. Also consider performing backups remotely from servers or workstations that do not carry a significant role on the network.

Table 10-2 provides a summary of Travan cartridge information.

Table 10-2 Travan Standards*

Cartridge Name and Alias	Native Capacity Compressed Capacity	Read/Write Compatibility	Read Compatibility
Travan-1 (TR-1)	400 MB 800 MB	QIC-80, QW-5122	QIC-40
Travan-2 (TR-2)**	800 MB 1.6 GB		
Travan-3 (TR-3)	1.6 GB 3.2 GB	TR-2, QIC-3020, QIC-3010, QW-3020XLW, QW-3020XLW	QIC-80, QW-5122, TR-1
Travan-8 GB (Travan 4, TR-4)	4 GB 8 GB	QIC-3095	QIC-3020, QIC-3010, QIC-80, QW-5122, TR-3, TR-1
Travan NS-8	4 GB 8 GB		QIC-3020, QIC-3010, QIC-80
Travan NS-20 (TR-5)	10 GB 20 GB	QIC-3220	Travan-8 GB, QIC-3095, TR-4

* Not every Travan standard is described in this book, but this chart will be useful when you need to determine backward compatibilities. Compatibilities may vary from one drive manufacturer to another.

** Travan-2 never really got off the ground.

Another type of tape drive found on backup systems is a **digital audio tape (DAT)** drive. DAT was originally designed as a high-fidelity digital replacement for standard analog audio cassettes and, as happened with CD audio disks, the format was quickly picked up by the computer industry. While the mechanics of reading and writing to Travan, QIC, and similar format tapes are analogous to the way audio signals are written to a standard audio cassette, writing data to a DAT is similar to the way video signals are written to a video cassette. Rather than the tape being moved linearly across a static head, it is moved across an angled, rotating head (see Figure 10-4). The result of this **helical scanning** is that a DAT can hold a higher density of data in a given area than other tape technologies. However, it is mechanically much more complex and potentially more expensive to repair should something go wrong. Such devices are usually significantly faster than more conventional linear drives both in the time taken to read or write data and in how quickly an individual file can be located.

Tape drives reading DAT pull the tape inside the drive to thread it. Be sure that you do not attempt to pull out a tape while it is active or you could have a real mess on your hands.

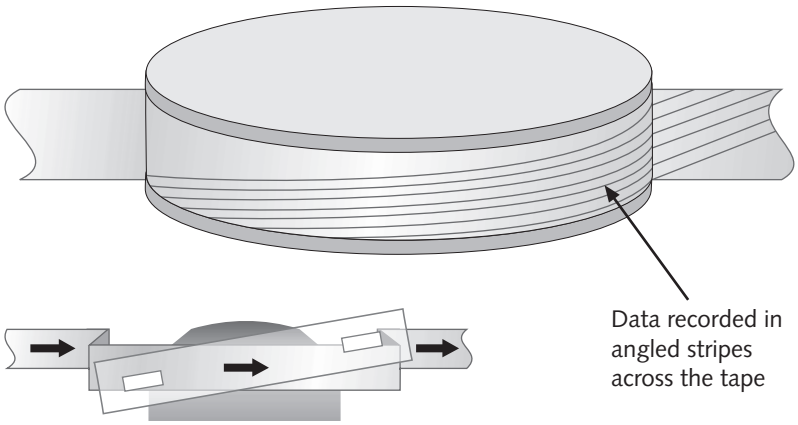

Figure 10-4 An angled DAT drive head stores more data on the tape

Only QIC tapes and drives read or write to QIC tapes and drives. In other words, you cannot use a digital tape in a QIC drive or a QIC tape in a digital drive. In the digital arena, larger tape drives of each category are backward compatible with smaller drives of the same type.

When Sony released its 8 mm video cassette technology, which is mechanically and conceptually similar to 4 mm DAT technology but with an 8 mm rather than a 4 mm tape width, its potential as a backup medium was once again seized upon by the computer industry. But unlike the case for 4 mm data drives, where a large number of makes and models are available, only a very few manufacturers took up the 8 mm helical scan challenge.

A variation of the 8 mm helical scan technology called **Advanced Intelligent Tape (AIT)** has recently become available (see Figure 10-5). AIT is a Sony invention that makes backup and restore operations faster because of an optional Memory in Cassette (MIC) chip that is able to quickly locate which of the 256 tape partitions contain the data. Compare this to traditional methods in which you must scan the tape to locate data. Larger amounts of compressed data can fit onto one tape, making it the highest-capacity helical scan format available and allowing it to compare more favorably in terms of capacity with another commonly found tape format, **digital linear tape (DLT)** (Figure 10-6).

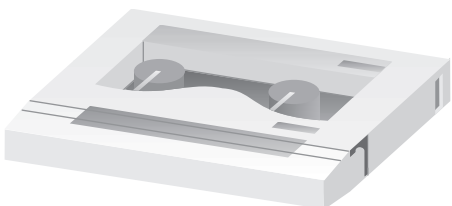

Figure 10-5 The AIT tape has a memory chip to quickly locate data

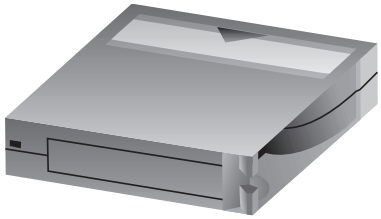

Figure 10-6 A DLT cartridge

Depending on the drive and media used, the DLT format allows up to 70 GB of compressed data to be stored on one rather large tape, which, unlike 8 mm or 4 mm helical scan technology, passes linearly over a fixed head.

Like Travan NS technology, DLT drives can simultaneously read and write, allowing them to perform extremely well, and in some cases even better than an 8 mm helical scan tape. Because they are in general significantly faster and more capacious than any other technologies, AIT, conventional 8 mm helical scan, and DLT drives are also more expensive.

One of the latest tape formats is the ultra-high-capacity **Ultrium** format. The Ultrium format is actually a subset of the **Linear Tape Open (LTO)** technology—a collaborative effort headed up by HP, IBM, and Seagate. Ultrium features include:

- *Single reel*—Internal cartridge mechanics maintain the tape, and because it is pre-threaded inside the cartridge (as opposed to threading tape through the tape drive), Ultrium-compatible drives will potentially be less complicated. Figure 10-7 shows a cutaway of the Ultrium tape cartridge.
- *High storage capacity*—Ultrium tapes offer a native capacity of up to 200 GB and data transfer of 20–40 MBps. The Ultrium LTO format is expected to be available in four different generations as advances in the technology develop, culminating in up to 1.6 TB per cartridge and up to 320 MBps (see Table 10-3).
- *Cartridge memory*—Ultrium cartridges can contain **LTO-CM (Linear Tape Open - Cartridge Memory)** right on the cartridge that stores a redundant file log and user-defined information. If you want to know what is on a given tape, you can use an external reader to access the memory. The LTO-CM uses an RF (radio frequency) interface. Compare this to the lengthy process required of other tape types for which you insert the tape into the drive and then use the backup software to build a catalog. This process can take several minutes, whereas Ultrium memory allows immediate access.
- *Error correction*—LTO technology provides two levels of error correction that can recover data even when longitudinal scratches appear on the media. Simultaneous read/write capability allows for real-time verification of data.

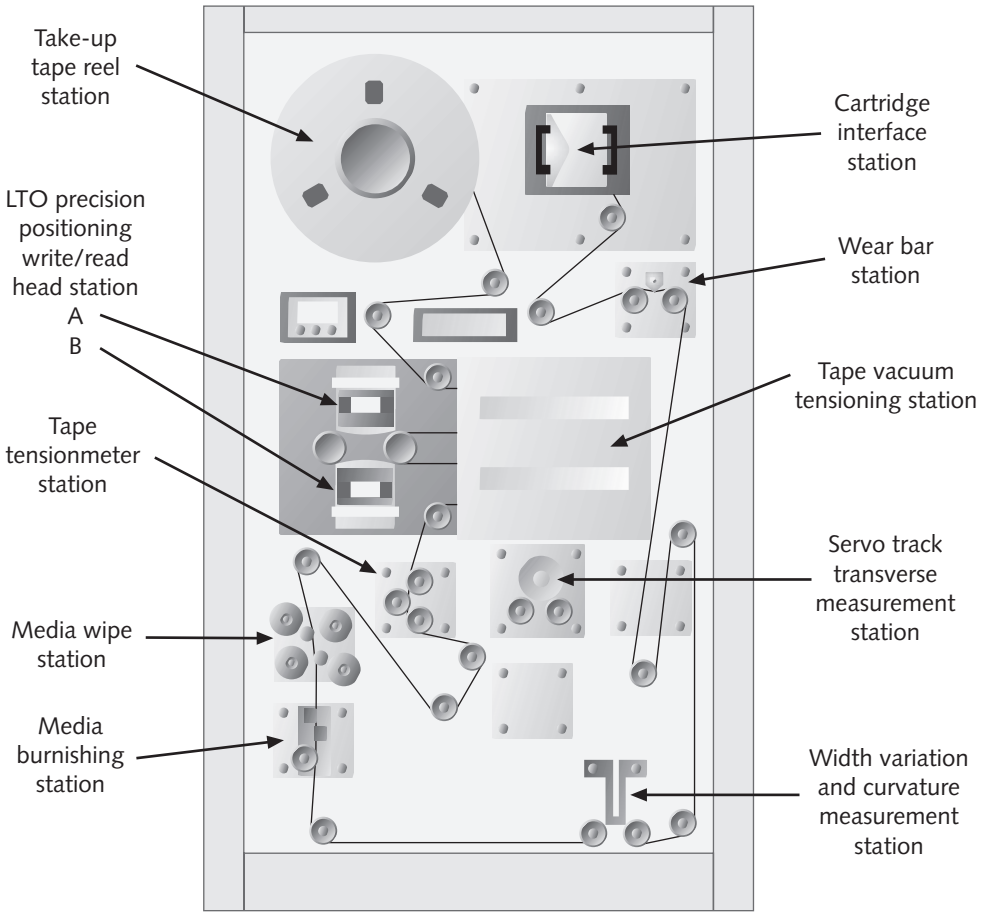

Figure 10-7 The Ultrium tape cartridge uses a single spool and is pre-threaded inside the cartridge

Table 10-3 summarizes information about the Ultrium format.

Table 10-3 Ultrium Format

Feature	Generation One	Generation Two	Generation Three	Generation Four
Capacity:				
Native	100 GB	200 GB	400 GB	800 GB
Compressed	200 GB	400 GB	800 GB	1.6 TB
Transfer Rate:				
Native	10–20 MBps	20–40 MBps	40–80 MBps	80–160 MBps
Compressed	20–40 MBps	40–80 MBps	80–160 MBps	160–320 MBps

Another LTO format known as Accelis proposes extremely fast data retrieval from anywhere on the tape in between 6.3 and 9.6 seconds. At this writing, Accelis is only a paper standard, and many speculate that Accelis drives and tapes may never actually be produced.

Tape and Tape Drive Maintenance

Tape media are relatively delicate, and you should exercise care when handling them. The magnetic metal oxide that coats the film on the tape is susceptible to wearing off slightly over time. For the most part, this won't have an effect on the ability of the tape to store data. However, it does affect the drive itself as literally miles of tape stream over the heads, capstans, and roller components. Just like a VCR, a tape drive requires regular cleaning of the components that come into contact with the media. Without proper cleaning, tape backups can lose integrity when the heads have difficulty reading or writing through the "gunk" that forms over the heads. Some newer media such as AIT include a built-in head cleaner; however, even in the cleanest environments tape drives will eventually accumulate contaminants. Dirty capstans and rollers can cause media to stick to the components and create a horrible mess when you have to untangle yards of loose tape from inside the drive. Good luck preserving the integrity of the tape after that! Some devices include automatic cleaning capability; otherwise, you will have to either procure cleaning tapes or use the old cotton swab and cleaning solution method.

Be careful about how you repurpose a previously recorded tape. In the past, when administrators had an old tape that they wanted to recycle to store new data, they would **bulk erase** the tape using a big magnet. Obviously, since tape is a magnetic medium, bulk erasing wipes out the existing data. With current tape technology, do *not* erase tapes in this way. Tapes now come preformatted, and bulk erasure will remove important markings from the tape. Once this occurs, there is no way for the tape drive to orient the tape to a known starting position or locate boundaries to logically store the data. If you want to erase a tape, use tape backup software, which can usually perform a quick erasure by removing the table of contents, or a secure erasure, which overwrites the tape with zeros and ones.

A tape written under one software vendor's program usually cannot be read by another vendor's program. This makes a corporate-wide policy that uses the same backup software critically important.

If you've performed even a single tape backup, then you know by listening that the drive fast-forwards and rewinds the tape quite a lot. This is okay, and is the nature of the backup. However, over time and especially if the tape does not go all the way from one end of the tape to the other, varying levels of tension will occur at various locations on the spool, and could affect read/write reliability. Temperature changes, which cause the tape to expand and contract, and dropping a tape can also affect its tension level. Tape

software offers utilities to **re-tension** the tape (see Figure 10-8), which fast-forwards to the end without reading or writing data, and then rewinds all the way to the beginning again. This process makes tension even throughout the tape.

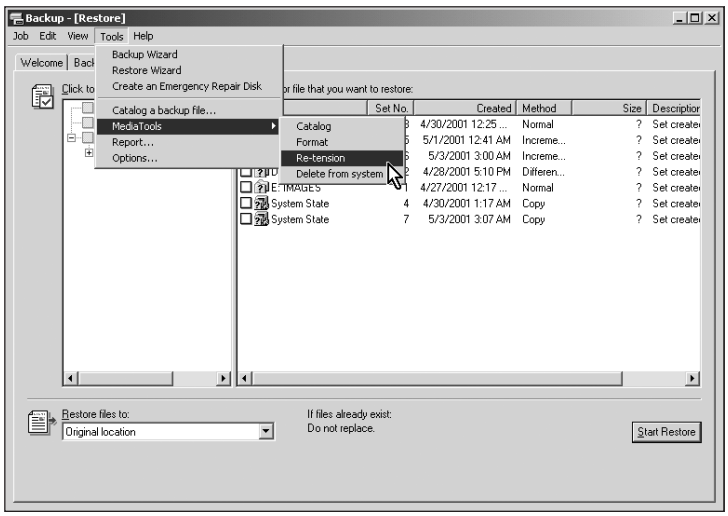

Figure 10-8 For even tension throughout the tape, the Windows 2000 Backup utility re-tensions the tape

If you need to retrieve data from a tape that has been erased or partially overwritten with other data, contact a data recovery specialist. Their services are very expensive, but they can retrieve data from tapes as well as failed hard drives, often with a high success rate.

Backup Software

On the whole, lower-cost tape drives, and some disk-based devices intended for use with stand-alone PCs, come bundled with backup software. The software tends to be quite basic, but it usually does the job it is intended to do. Such software is not capable of servicing more sophisticated backup solutions such as automatic tape changers, remote backup, or backing up open files. On the other hand, most high-end tape drives and most disk-based products do not come bundled with software, and this should be taken into consideration when comparing costs, because the software is an additional several hundred dollars per server. Software for high-end tape drives includes many more features, and at a reduced cost per remote station.

If you need software as part of your backup setup, there are a number of packages available, ranging from replacements for bundled software to high-end products capable of backing up a whole room full of servers onto several tape drives, in some cases

simultaneously. Depending on your needs, features to look for in backup software include the following:

- The ability to restore your hard disk in its entirety without requiring you to reinstall the operating system manually. This software feature creates an image of the disk and stores it on tape. This is not unlike Symantec Ghost or PowerQuest DriveImage, which store the image on a hard drive.
- The ability to treat a backup tape as a virtual hard disk, albeit a very slow one. This capability sometimes uses the right-click menu from Windows Explorer, allowing you to perform the backup in a similar way to using the Send To item when copying a file to a floppy disk.
- A built-in data compression capability should your drive not include its own hardware data compression facilities.
- The ability to scan for viruses during backup or restore operations.
- Built-in tape management facilities that can tell you when to replace or swap tapes. Robotic or automated media-switching devices will often require proprietary drivers and administrative software to be installed.
- The ability to create an account for you that has access rights to back up all data regardless of ownership. In a Windows NT/2000 environment, this account would belong to the Backup Operators group. The account is usually able to log on as a service and, therefore, perform backup without an administrator performing an actual logon.
- The ability to perform unattended backups. The software often creates a special account that can log on by itself at scheduled times to start the backup.
- The ability to run commands before and after the backup. This is useful for stopping services or applications that would interfere with the backup utility and restarting them afterward.

Storing and Securing Backups

There are several steps that you can take to enhance the security and operation of your backup and restore operations. You should also take steps to secure your backup cartridges.

When you develop a backup plan, consider the following recommendations:

- Secure both the storage device and the backup cartridges. Data can be retrieved from stolen cartridges and restored to another computer.
- If the software supports it, add security measures to the tape. You might be able to protect the media with a password or allow restore operations only by the Administrator or Owner accounts as shown in the Windows 2000 backup in Figure 10-9.

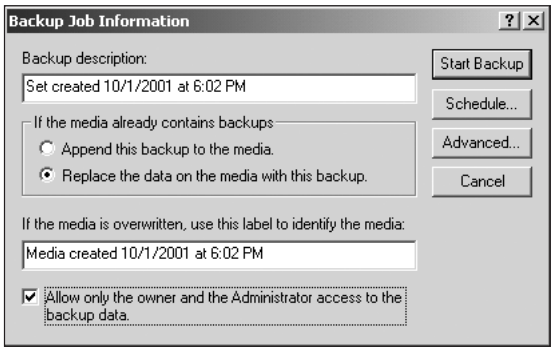

Figure 10-9 You can specify that only the Owner or Administrator accounts are allowed to restore backups

- Back up an entire volume by using the normal backup procedure. In case of a disk failure, it is more efficient to restore the entire volume in one operation than to also include differential or incremental backups.
- Keep at least three current copies of backup cartridges. Store one copy at an off-site location in an environmentally controlled, secure environment. Check for a service bureau in your area that can provide this storage. Most service bureaus also have tape drives that can read and restore the tapes for you if necessary. In preparation for disaster recovery, know how long it takes to physically retrieve the tapes from off-site locations. Store another copy near a server that can restore from the tape in a secure, locked, fireproof cabinet. The last tape can be stored wherever it suits you; its main purpose is redundancy in case one of the other tapes becomes damaged or defective. Some organizations with a WAN infrastructure send a copy of the normal backup to another office across the WAN. That way, if the local off-site location and the local office become unavailable (as in a natural disaster), you still have the copy that was sent to the other office.

Backup Types

A **normal backup** copies all selected files and clears the archive bit on each one. This identifies the file as having been backed up. The next time the file is modified, the archive bit is automatically set (added). Files or directories that have been moved to new locations are not marked for backup. Most backup software allows you to back up only files with this marker set and to choose whether or not to mark files when they are backed up. The archive bit is significant in relation to the incremental and differential backups. Normal backups are the easiest to use for restoring files because you need only the most recent backup file or tape to restore all of the backed-up files. Normal backups take the most time because every file that is selected is backed up, regardless of whether it has changed since the last backup.

An **incremental backup** backs up only those files that have been created or changed since the last normal or incremental backup, which can reduce the amount of time that is required to complete the backup process. It marks files as having been backed up by clearing the archive bit. You should create a complete normal backup of your system before you run incremental backups. If you use a combination of normal and incremental backups, you must have the last normal backup set as well as every incremental backup set that has been made since the last normal backup—in chronological order—to restore your data.

The archive bit is easily observable by viewing the properties of a file. For example, in Windows Explorer, you would right-click a file and view its properties.

The advantage of an incremental backup is that the backup process is typically faster than both a normal or differential backup. The disadvantage is that it takes longer to restore the backup because you might have to supply multiple incremental tapes created since the last normal backup.

A **differential backup** copies files that have been created or changed since the last normal or incremental backup, which can reduce the amount of time that is required to complete the backup process. It does not mark files as having been backed up. You should create a complete normal backup of your system before you run differential backups. If you are doing normal and differential backups, you must have the last normal backup set and the last differential backup sets to restore your data.

The advantage of the differential backup is that it is faster than the normal backup and only requires two backup sets to restore data: the original normal backup and the corresponding differential backup. The disadvantage is that differential backups take longer to back up data.

A **copy backup** copies all selected files, but it does not mark each file as having been backed up. Copying is useful to back up files between normal and incremental backups because it does not affect other backup operations. A **daily backup** copies all selected files that have been modified on the day that the daily backup is performed. The backed-up files are not marked as having been backed up.

Backups protect against data loss caused by a virus. Because some viruses take weeks to appear, keep normal backup tapes for at least a month to make sure that you can restore a system to its uninfected status.

Using Incremental Backups

Let's assume that you implement a normal backup, and that you do not want to do a normal backup every day. Performing an incremental backup will make the most sense in order to keep track of what files are backed up (unlike a differential backup). For

example, suppose you're doing an incremental backup of 850 MB of data. If you did the normal backup on Monday, all 850 MB would be placed on the tape on that day. On Tuesday, the incremental backup day, only the files that you've created or changed since Monday would be backed up—so you may back up only a few megabytes at that time. On Wednesday, you'd back up only the files that you've created or changed since Tuesday, and so on. Because each incremental backup only backs up data since the last normal or incremental backup, the backup is relatively fast.

Unfortunately, incremental backups have their disadvantages. Although they are quicker than normal backups, it takes longer to restore a file when incremental backups are involved in the process. For example, suppose you needed to restore a file you backed up on Monday. This file would be included in the first backup of the set, since you performed a normal backup on Monday. However, to ensure that you were restoring the most recent version of the file, you'd have to search all remaining backups in the set to see if the file had changed since then. Another disadvantage to this type of backup is the fact that you must have all backups in the set available in order to restore only one file.

If you're planning to use an incremental backup, you might consider starting a new backup set every Monday that would include the backups for Monday through Friday of that week. Of course, if your servers are heavily utilized on weekends, you can schedule Saturday and Sunday backups as well. You can also make the length of time between creating new backup sets as long or as short as your needs dictate.

Develop a Backup Strategy

Regular backups of local hard disks prevent data loss resulting from disk drive failures, disk controller errors, power outages, virus infection, and other possible problems. Backup operations that are based on careful planning and reliable equipment make file recovery easier and less time consuming. There are several backup strategies, though the most common strategy is known as the Grandfather-Father-Son (GFS) strategy (described below). Your backup strategy will usually use some combination of normal, incremental, or differential backups.

Developing a backup strategy involves not only determining when to perform backups, but also testing the data with random and scheduled verification to make sure that tape devices and media are functioning properly. Also, make sure that throughout the organization, the persons or departments responsible for handling backup and restore operations are well informed as to what exactly they should back up. You would not want to be caught in the awkward situation where someone asks you to restore a file that you didn't even know you were supposed to back up.

The Grandfather-Father-Son Backup Strategy

There are several generally accepted backup strategies; however, many of them are based on the popular **Grandfather-Father-Son (GFS)** backup strategy (otherwise known as

the Child-Parent-Grandparent method). This backup strategy uses three sets of tapes for daily, weekly, and monthly backup sets, and you implement it as follows:

1. *Back up the "Son"*—Label four tapes as "Monday" through "Thursday." These Son tapes are used for daily incremental backups during the week. For subsequent weeks, reuse these same tapes.
2. *Back up the "Father"*—Label five tapes as "Week 1" through "Week 5." These Father tapes are used for weekly normal backups on Friday, the day you do not perform a Son backup. Once you make the tape, store it locally. Reuse the tapes when each tape's respective week arrives. Depending on your backup policy, periodically duplicate a Father tape for off-site storage. You can use another drive to perform a simultaneous backup, or some backup software might offer a tape copy feature.
3. *Back up the "Grandfather"*—Grandfather tapes are used for a normal backup performed on the last business day of the month. No standard labeling scheme is stated, but consider labeling three tapes as "Month 1" through "Month 3." The tapes are valid for three months and are reused every quarter.

At a minimum, the GFS strategy requires 12 tapes if you add them all together, assuming that no one backup exceeds the capacity of a single cartridge.

Of course, you can modify this scheme as it suits your backup policy, but the GFS strategy is a logical, reliable place to start. For example, if you want to keep a year's data archived at all times (instead of only a quarter's), then for the Grandfather tapes you would label 12 tapes "Month 1" through "Month 12" and reuse the tapes every year. An illustration of the GFS rotation scheme appears in Figure 10-10.

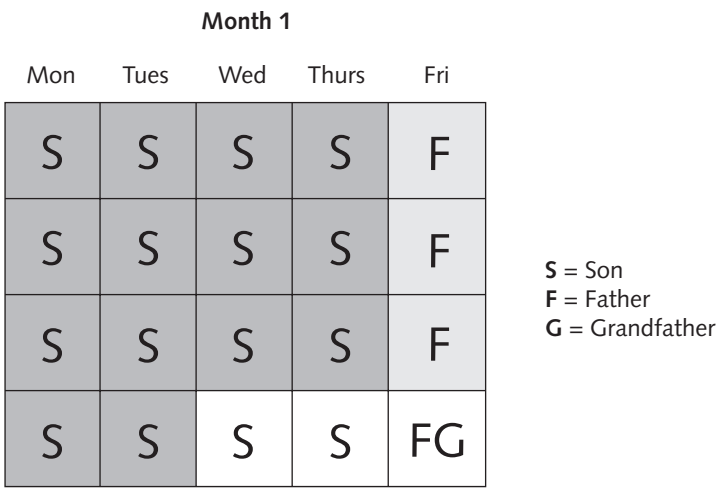

Figure 10-10 The GFS tape rotation strategy

The Six-Cartridge Backup

If you want to use fewer tapes, consider the **six-cartridge backup** strategy. This might be a better choice for smaller business or sites that do not generate large quantities of data. The disadvantage of the six-cartridge strategy is that you have a shorter archive history—only two weeks, whereas the GFS strategy is a quarter.

To perform a six-cartridge backup:

1. Label six cartridges "Friday 1," "Friday 2," "Monday," "Tuesday," "Wednesday," and "Thursday."
2. Perform the first normal backup onto the "Friday 1" tape. Store the tape off-site.
3. On Monday, perform an incremental backup onto the "Monday" tape. Store the tape on-site.
4. Repeat the incremental backup onto tapes "Tuesday," "Wednesday," and "Thursday" on the corresponding days.
5. Perform the second normal backup onto the "Friday 2" tape. This completes the backup cycle. Store the cartridge off-site. On each successive Friday, alternate between the "Friday 1" tape and the "Friday 2" tape.

Note that the Friday tapes are always stored off-site, as illustrated in Figure 10-11. This can cause a problem when it is necessary to restore data, because you will have to physically retrieve the tape, and this might not suit the time frame necessary to restore the data. As with the GFS method, you can modify the six-cartridge method to suit your needs.

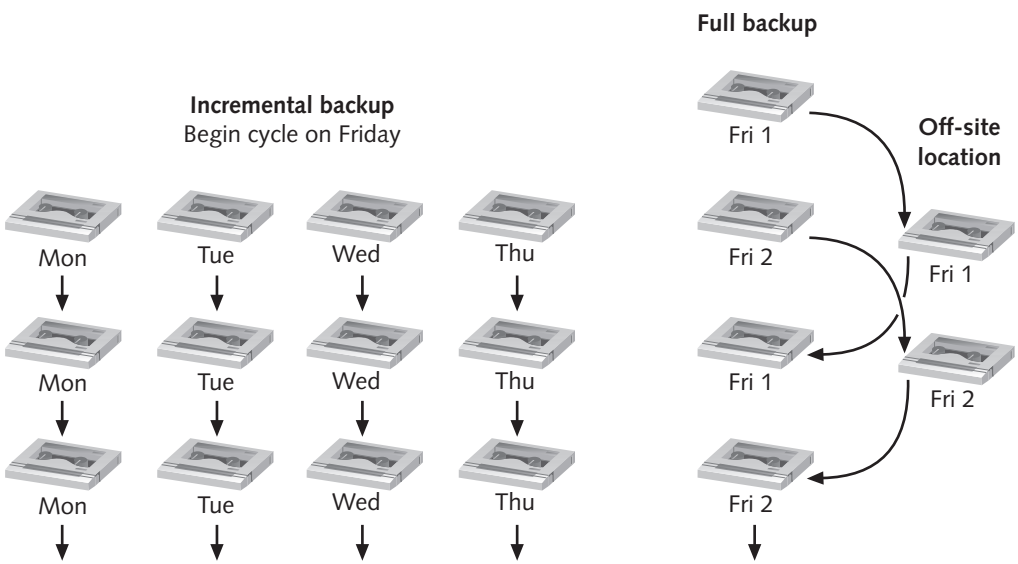

Figure 10-11 The six-cartridge backup strategy stores Friday tapes off-site

To avoid the need to retrieve a Friday tape from off-site storage when you need to restore data, make the Monday backup a normal backup and the Tuesday, Wednesday, and Thursday backups incremental to the Monday backup.

The Tower of Hanoi

Finally, there is the **Tower of Hanoi backup** method, borrowing from a mathematical logic game of the same name. Use five sets of media for this rotation, labeling each "A" through "E" and proceed as follows:

1. On day one, back up to "A." Reuse the "A" tape every other day.
2. On day two, back up to "B." Reuse the "B" tape every four days.
3. On day four, back up to "C." Reuse the "C" tape every eight days.
4. On day eight, back up to "D." Reuse the "D" tape every 16 days.
5. On day 16, back up to "E." Reuse the "E" tape every 32 days.

As a memory aid to the Tower of Hanoi rotation, just notice that each tape is reused in a pattern similar to binary notation. Look at the "reuse" schedule at the end of every step above. Notice that you reuse tapes every 2, 4, 8, 16, and 32 days.

An advantage of the Tower of Hanoi rotation is that you always have a daily history of data extending back 32 days. This is a flexible backup strategy requiring only five tapes (assuming each backup only requires one tape). If you want long-term archiving, you can remove a tape and store it off-site. Label the tape with the date and replace it with another. For example, if you want to archive the "E" tape, place a date on it, send it to storage, and label a new cartridge "E" that will continue the rotation. Also, you can extend the history if you like. By adding an "F" backup every 64 days, you now have a 64-day history. Keep adding letters until you reach the history you want.

The rotation scheme is best understood by viewing Figure 10-12.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
A		A		A		A		A		A		A		A	
	B				B				B				B		
			C								C				
							D								
															E

return to day 1

Figure 10-12 Tower of Hanoi media rotation schedule

Document the Process

Keeping accurate backup records is essential for locating backed-up data quickly, particularly if you have accumulated a large number of backup cartridges. Thorough records include cartridge labels, catalogs, and online log files and log books.

- *Cartridge labels*—Cartridge labels for write-once cartridges should contain the backup date, the type of backup (normal, incremental, or differential), and a list of contents. If you are restoring from differential or incremental backups, you need to be able to locate the last normal backup and either the last differential backup or all incremental backups that have been created since the last normal backup. Label reusable media, such as tapes or removable disks, sequentially, and keep a log book in which you note the content of cartridges, the backup date, the type of backup, and the date the medium was placed in service. If you have to replace a defective cartridge, label it with the next unused sequential ID, and record it in the log book.
- *Catalogs*—Most backup software includes a mechanism for cataloging backup files. Backup software typically stores backup catalogs on the cartridge and temporarily loads them into memory. Catalogs are created for each backup set or for each collection of backed-up files from one drive.
- *Log files*—Log files include the names of all backed-up and restored files and directories. A log file is useful when you are restoring data because you can print or read this file from any text editor. Keeping printed logs in a notebook makes it easier to locate specific files. For example, if the tape that contains the catalog of the backup set is corrupted, you can use the printed logs to locate a file.

High Server Availability and Redundancy

File services, print services, and client-server applications rely not only on the availability of the computers running the services, but also on the availability of network services. In an environment where there is only one computer providing a particular service (file, print, or application), an outage involving that server eliminates the availability of the provided service. To provide both load balancing and redundancy for these services, a group of computers (a cluster), can cooperate in providing these services. This cooperation is managed by clustering software that provides a service to clients in a client-server environment. For example, a public file share, a web server, or a database application can all be managed as resources.

A cluster improves the availability of client-server applications by increasing the availability of server resources. Using a cluster, you can set up applications on two or more servers (nodes) in a cluster. Each node connects to a shared storage media. Clusters present a single, virtual image of the cluster to clients (see Figure 10-13). If one node fails, the applications on the failed node are available on the other node. Throughout this process, client communications with applications usually continue with little or no interruption. In most cases, the interruption in service is detected in about five seconds, and services can be available again in as few as 30 seconds (depending on how long it takes to restart the application).

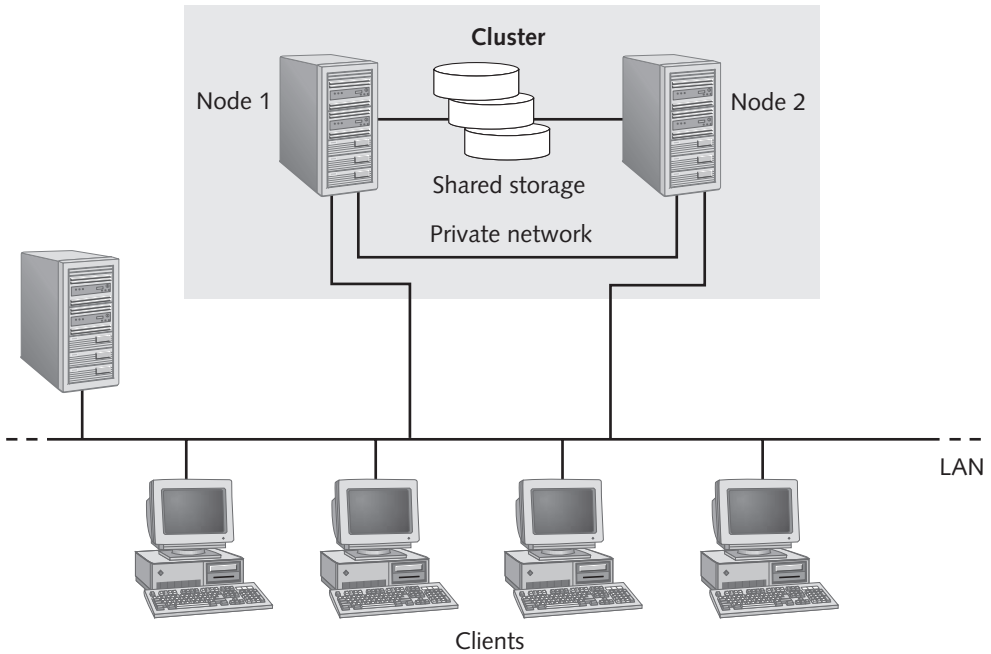

Figure 10-13 Cluster technology connects two or more servers to common shared storage

Clustering provides high availability and fault tolerance by keeping a backup of the primary system available. **Static load balancing** remains idle and unused until a failure occurs, which makes this an expensive solution. An **active cluster** is a clustering method in which all nodes perform normal, active functions and then perform additional functions for a failed cluster member. For example, redundant systems might have one node in the cluster servicing web clients while the other node provides access to a database. If either node fails, the resource (either the web server or database server) fails over to the other node. The node that is still functioning responds to both web and SQL requests from clients (see Figure 10-14). In a **passive cluster**, a server with identical services as its failover partner would remain in an idle node state until such a time as the primary node fails.

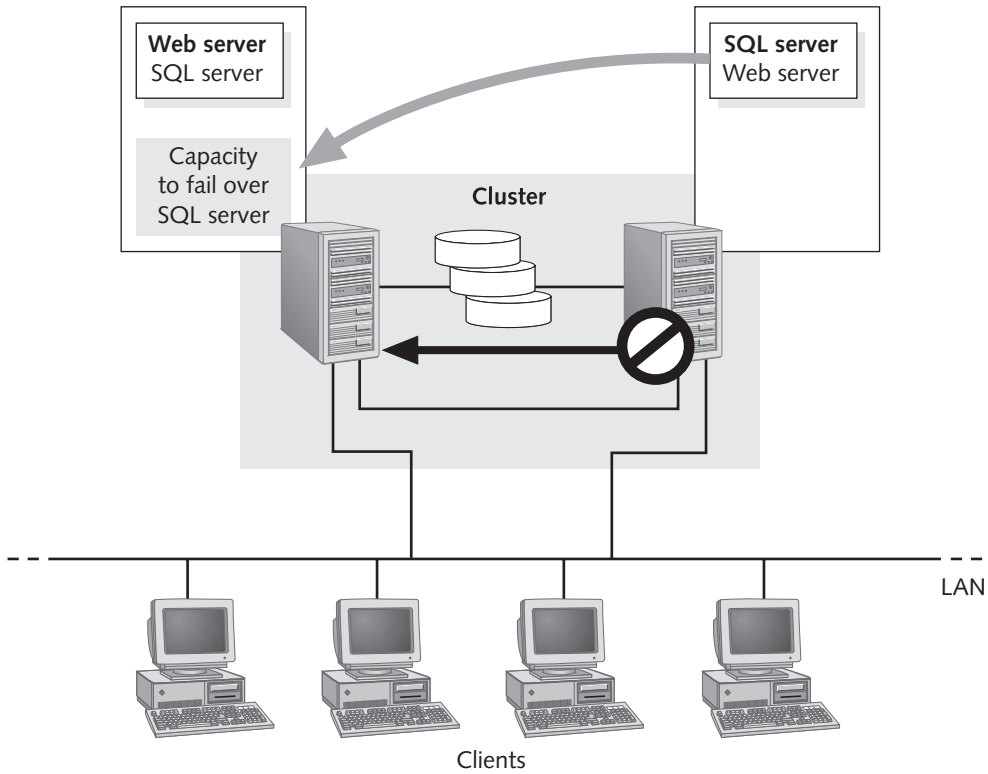

Figure 10-14 Active cluster redundancy

CPU, I/O, storage, and application resources can be added incrementally to efficiently expand capacity, making the solution highly scalable. This creates reliable access to system resources and data, as well as investment protection of both hardware and software resources. Clusters are relatively affordable because they can be built with commodity hardware (high-volume components that are relatively inexpensive).

By clustering existing hardware with new computers, you protect your investment in both hardware and software: Instead of replacing an existing computer with a new one of twice the capacity, you can simply add another computer of equal capacity. For example, if performance degrades because of an increase in the number of clients using an application on a server, you can add a second server to a cluster, which improves performance and also increases availability.

Clients typically access network applications and resources through network names and IP addresses. When these network applications and resources are hosted within a Microsoft Cluster Server (MSCS), clients can continue to find and access the resources, even though they may move between nodes. MSCS enables this by failing over both the IP address and network name for a given resource.

Failover

Failover is the process of having cluster resources migrate from an unavailable node to an available node. A related process, **failback**, occurs when service transfers back to the node that has been offline after it is back online. The cluster automatically initiates failover when it detects a failure on one of the cluster nodes. Because each cluster node monitors both its own processes and the other node, the need for failover is detected without delay.

Spare Parts

For hardware failure recovery, having a number of spare parts available will save the time needed to order failed items from the original vendor. Also, as equipment ages, the availability of the parts needed to continue operation may diminish as well. All parts to be considered available as replacements for a given computer system must be compatible with both the operating system in use and other components within the system. The best strategy for mission-critical systems is to have a set of matching parts available. Some vendors, such as Intel, sell spare-parts kits comprised of the most critical system components, including:

- 12 V VRM (a voltage regulator module that helps ensure clean power to the motherboard)
- Fans
- Hot-swap bay assembly with SCSI backplane
- Power supply
- CD-ROM drive
- Floppy drive
- Cables

Using the spare-parts kit, you use parts as necessary and immediately call the vendor to replace the parts you use. That way, the spare-parts kit is always ready for service. Be sure to also have additional spare parts, which might not be included in a spare-parts kit, such as:

- Network card
- Memory modules
- Processor
- Hard disk
- Hard disk host adapter (EIDE or SCSI)
- Motherboard
- Video card
- Sound card (optional)
- Other miscellaneous I/O boards

Recall from earlier discussions that hot swapping will allow you to replace many of these devices without an interruption in service. This is particularly true of hard disks. Many systems also include a **hot spare** that is connected to the hard disk host adapter along with the other hard disks, but is dormant until another device in the drive array fails. At that time, the hot spare is usually placed into service automatically. For example, if one member of a mirrored (RAID-1) array fails, the hot spare can automatically come online and the remaining member will begin to duplicate to it. The main difference between a hot-swappable disk and a hot spare is that a hot spare is not a Plug and Play item. It must be on the bus at the time of the failure. If it is not, then you must shut down the system and add it to the bus, which would not be necessary with a hot-swappable drive.

Be sure that you secure the spare parts in a locked cabinet. Most of the parts are small and easily stolen.

SNMP Settings

In Chapter 9, you learned the basics of SNMP. SNMP is a critical element for disaster avoidance, detection, and recovery. The following is a short list of common SNMP items for which you will want to configure your SNMP management system to assist in disaster prevention and recovery:

- Network protocol identification and statistics
- Dynamic identification of devices attached to the network (discovery)
- Hardware and software configuration data

- Device performance and usage statistics
- Device error and event messages
- Program and application usage statistics

The items in the list can help to detect impending problems and verify that a proposed solution worked effectively. For example, if you suspect network traffic to be excessive on an Ethernet network, you could configure the SNMP agent to issue a trap when it detects over 30% network utilization on a given segment. After replacing a hub with a switch to increase throughput and reduce collisions, the same SNMP agent can confirm (by absence of a trap) that the solution worked.

In general, agents do not originate messages—they only respond to messages. The exception is an SNMP trap triggered by a specific event such as a system reboot or illegal access. Traps and trap messages provide a rudimentary form of security by notifying the management system any time such an event occurs. Typically, you configure the management system to issue an email, fax, network message, or pager alert. For pager alerts, some systems require an external modem, but I usually circumvent that by using email instead to send text messages to a cell phone or email-capable pager.

If your cabinet and SNMP management software support it, you can add a layer of physical security to the server cabinet by configuring SNMP to issue a trap any time the cabinet door opens. If you're the only one who is supposed to access the cabinet and your pager receives an SNMP message while you're away at lunch, then you know somebody is illegitimately accessing the cabinet.

Server Management and Maintenance

Server management and disaster recovery are really two balancing components in the same overall server health management scheme. You use server management software and faithful physical management of the server to prevent disaster. Then, you use disaster recovery techniques to fix the inevitable problems that occur.

Server Management Software

There is a broad selection of server management software. In earlier chapters, we have mentioned third-party management products such as IBM Tivoli or Computer Associates' Unicenter TNG. In this section, third-party products can still play a role, but servers also usually include less comprehensive software that provides basic system monitoring functions. These utilities often integrate with the system BIOS or CMOS settings, and display or issue an alert when a problem appears. For example, Figure 10-15 shows a very simple server monitoring utility that monitors temperature, fan speed, voltages, and more.

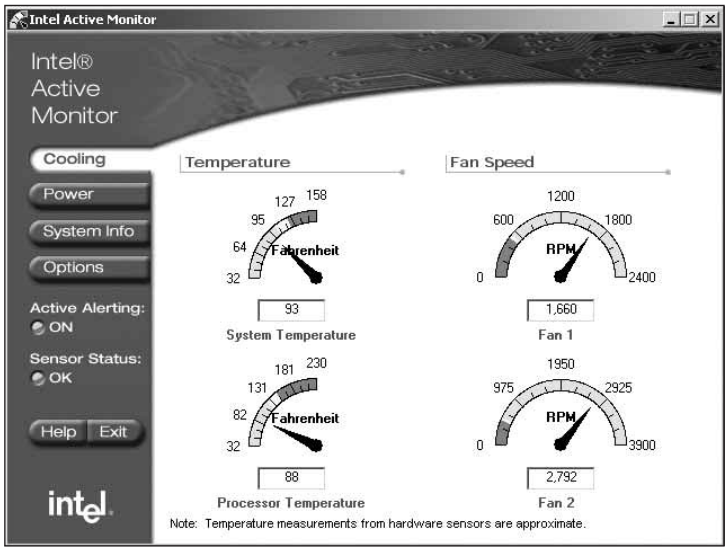

Figure 10-15 A monitoring utility displays basic server health issues

10

Server motherboards and management software usually offer features such as:

- *Failure detection*. Detects changes in temperature, voltage, fan speed or failure, disk drive problems or failure, power supply failure, processor status, and ECC memory errors.
- *Software monitoring*. Detects hung applications. For serious problems, you can use management software to perform a graceful shutdown or reboot.
- *Event logging*. Events are stored in NVRAM (nonvolatile RAM) so that if power is lost, the records remain.
- *Emergency Management Port* (Intel boards). A feature that allows you to remotely turn on, off, or reset the server and view the event log. These features require an external modem and are very useful for remotely monitoring servers over a wide geographic area from a single location.
- *Security monitoring*. A jumper setting enables chassis intrusion detection. Some systems will automatically blank the video when the chassis or cabinet is open, and a password is required to resume normal video.

Most server boards include a server management utility, but if yours does not, I recommend downloading at least a monitoring utility such as the Motherboard Monitor (freeware) from nearly any popular download site, or *www.tweakfiles.com/diagnostic/motherboardmonitor.html*.

Larger organizations need more than a local server management utility, and opt instead for enterprise management software such as HP's Openview or CA's Unicenter TNG.

A very significant advantage of this type of software is that you can manage the entire network from a single seat, regardless of the physical location of any one server. For example, in Figure 10-16, I used CA's Unicenter TNG to access the California network and, in Figure 10-17, the North American network. The figures are shown using a map feature of Unicenter in 2D, but you can also use 3D graphics. Using the map view, you click on any of your networks to access information and administer it.

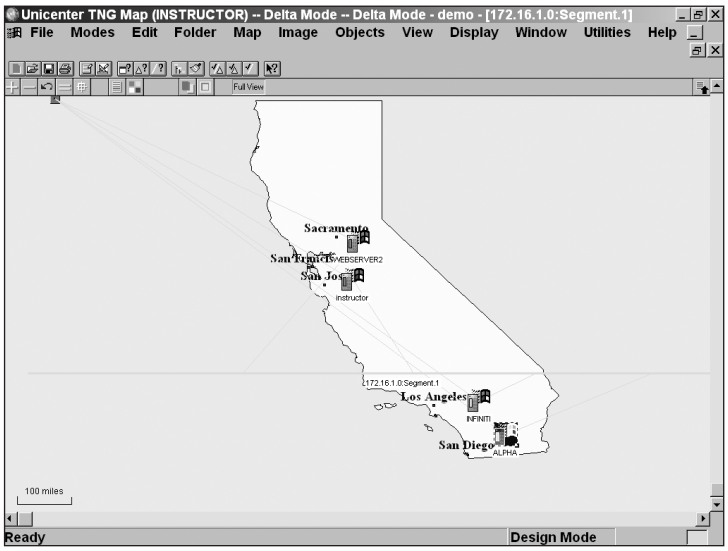

Figure 10-16 Unicenter TNG allows you to visually access your server

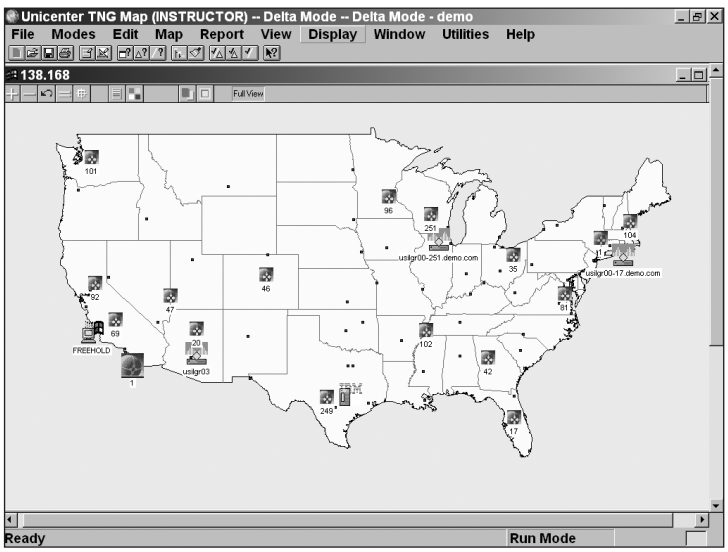

Figure 10-17 A demonstration of the Unicenter TNG map feature

Physical Care and Maintenance

Recall from Chapter 2 that two primary factors contribute to hardware device and computer peripheral failure: dirt and heat. Regular cleaning of computer equipment and adequate ventilation are necessary in order to maximize the lifetime of the equipment.

Hard disk drives, for example, are prone to failure in high-heat environments. Their mechanical nature causes a great deal of friction, as the platters can spin in excess of 15,000 revolutions per minute (rpm). Stack several disks inside of a single computer without proper ventilation, and the combined heat of several drives can damage the electrical components, leading to drive failure. If a drive fails and data is lost prior to a timely backup, an expensive data recovery service bureau will need to be employed to open the disks and attempt to extract the data from the failed device.

Consider checking into a server chassis that includes airflow guides. These guides are paddles near the cooling fans that allow you to direct the airflow as desired. This is very handy when the drives or other hot components are located some distance from the actual fans or are not in the path of the normal airflow. For example, Supermicro offers airflow guides on some of its popular models (see Figure 10-18).

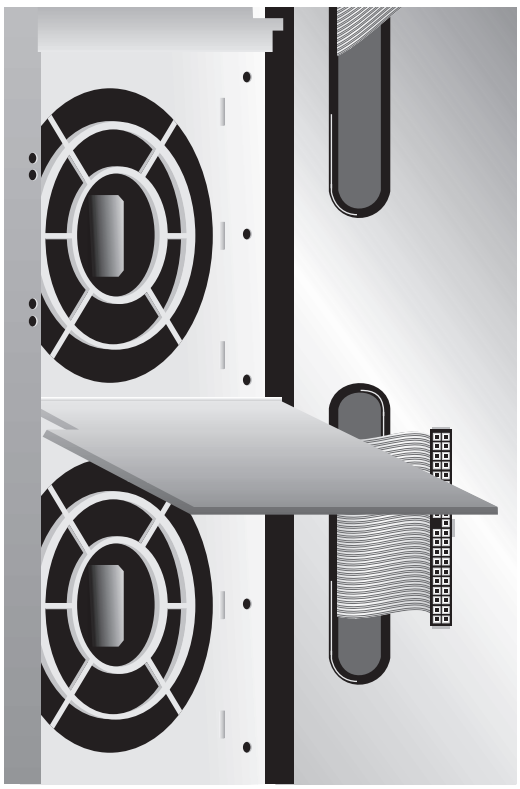

Figure 10-18 You can direct the airflow using adjustable airflow guides

Related to the heat and dust that accumulate during the normal use and function of your equipment is the age of the equipment and/or the time that this equipment has been in service. Regular monitoring of your system's performance is necessary to note changes in I/O and other performance measures unrelated to changes in network access or activity.

If a device is performing slowly or is experiencing an increase of random or unexplained errors, it may be in an early stage of failure and should be replaced as soon as possible. Some devices, such as a power supply, cannot be easily monitored without server management software, but can lead to systemwide failures if performance begins to deteriorate. A power supply whose voltage is beginning to drop significantly or fluctuate can affect many components in the system both in terms of quality and length of service. If you suspect a power supply failure, replace it immediately to avoid further damage to other system components.

Probably the most common maintenance issue for servers is planned downtime to blow out dust. The frequency of this task varies greatly depending on the environmental conditions at the site. Until you can determine the rate at which dust accumulates in the case, check accumulation weekly. Once you determine an optimum frequency for this planned downtime, be sure to schedule it regularly. When dusting, remember to remove the server from the server room if you are blowing out dust in order to avoid contaminating the environment. Remove dust from every place you can, though the chassis fans and power supply fans accumulate dust most quickly.

Don't forget to blow out CD-ROM and floppy drives. The floppy drive in particular is relatively open, even with the protective flap that covers the opening. I recently threw away about three floppies that the drive wouldn't read. I realized that it was too much of a coincidence for all of the disks to be bad. So I removed the floppy, blew out the dust, and almost choked on the cloud.

Developing a Disaster Recovery Plan

A **disaster recovery plan (DRP)** for a large enterprise can amount to hundreds of pages with thousands of contingencies. These complicated details and scenarios, if left unplanned, can go unaccomplished, leading to an unacceptable extension of system outage. When putting together a DRP, there are many items that need to be considered and included. Practical implementations will vary from one network to another, but the most important matter is to develop a disaster recovery plan in the first place.

The DRP is critical because in a disaster, the tendency is to hastily assemble a short list of recovery steps that come to mind. This can actually extend the time it takes to recover from disaster, because it is difficult to account for every recovery procedure. A DRP specifies what actions need to be taken, in what order, after the destructive event. Each scenario, from a single desktop computer failure to a complete outage of an entire site,

should have a plan of action for those responsible. Depending on the severity of the event, these actions can include:

- Evacuation of facility and notification to emergency services
- Notification sequence for team leaders and backups
- Establishing a temporary business recovery command center
- Preliminary and detailed damage assessment
- Recall of vital records from off-site storage

In the event of a site-wide or catastrophic failure, longer-term issues must be addressed:

- Handling legal, financial, and insurance issues
- Dealing with the news media to mitigate misinformation
- Locating interim facilities to restart your business
- Recovery of PCs, LANs and midrange systems
- Establishing voice and data communication
- Addressing human resource and accounts payable/receivable issues
- Replacement of equipment, furniture, and supplies
- Notification to clients, customers, suppliers, and stockholders

A disaster recovery plan is only a possibility when there is actually a team to develop it. Assemble members from each major branch or function of your organization, because disasters affect all aspects of the organization, not just the IT department. The disaster recovery team needs to work well together in order to minimize downtime and loss of productivity. Preparing a disaster recovery plan can take months and, depending on the size of the organization, perhaps over a year.

For an example of comprehensive disaster recovery services, see BMS Catastrophes (*www.bmscat.com*). For help in planning a disaster recovery solution, see Davis Logic (*www.davislogic.com*).

Using Alternate Sites

In larger computing environments, it is not practical to have enough computers in store and configured to replace most of the equipment all at once. Larger companies, universities, and institutions cannot afford to carry more that 1% or 2% of the total hardware collection in recovery-based inventory. Also, in the event of facility-, site-, or campus-wide failure, there may be more than computers that need to be replaced: Facilities may also have to be recovered as in the case of a flood, tornado, or other natural disaster.

In preparation for large-scale recovery, consider alternate sites. They can range from simple data collection and warehousing services designed to return data as needed to continue business operations, to full-scale facility duplication (**hot sites**) designed to replicate all of the hardware, software, and data infrastructure necessary to assume all business functions in the event of a primary facility failure. These hot sites are very expensive to maintain, as they have nearly a complete replica of all computers in the central facility, but if a large-scale outage will cost the company its existence, the price is worth it.

Hot Sites

Usually, a hot site is a shared facility with a number of subscribers from different geographic locations, each of which share in the cost of maintaining the fully operational center. Each subscriber usually occupies the hot site for up to six weeks after a disaster. These facilities are also available for subscribers to exercise their recovery plan in test mode.

Hot site recovery is appropriate for a computer operation that has more than a 24-hour outage tolerance. Data centers requiring faster service restoration must invest in redundant (spare) equipment that is immediately available to satisfy this need. Conversely, facilities that can afford to wait several weeks before restoring service need not engage a hot site, for they will have time to order and install new equipment.

Cold Sites

Many service bureaus augment their hot site service with a **cold site** feature. This is a facility designed to receive computer equipment. All power, water, air conditioning, raised floor, and other items requiring a long lead-time to acquire, install, and make ready to house a computer center are in place.

Should a company be unable to return to its home-computing center within a tolerable time frame, it would make arrangements to occupy this cold site. Computers, peripheral equipment, and related services would be ordered (purchased, leased, or rented) and made ready to assume the company's processing workload. Cold sites would be used until the home site was repaired or rebuilt.

Site Management

The business processes can be quite complex as remote sites, either hot or cold, are implemented. Bad planning in these areas leads to confusion and delay. Several important steps are necessary to determine the activities and timeframes surrounding service and equipment movement and delegation between the central (or original) site, hot sites, cold sites, and their return to the repaired or newly constructed site.

- Determine the extent of the damage and if additional equipment, services, and supplies are needed.
- Take care to adequately cover telecommunications issues. Most medium-sized and large organizations usually have a number of telecommunications services

(such as leased lines and fiber) connecting campuses and other facilities. This means that, in the event of a disaster, such connections to other facilities may have to be abandoned, re-routed, or installed to the hot and/or cold sites in order to establish connectivity to such other operational facilities. Detailed records of current operations must be carefully reviewed to be certain that your disaster recovery plan covers every conceivable technical aspect of recovering all critical services and data in the least amount of time.

- Obtain approval for expenditure of funds to bring in any needed equipment and supplies. I recently read of a corporation that set up an agreement with their bank so that in the event of a catastrophic disaster, the bank would supply a mobile branch staffed with at least two tellers who would dispense the finances and keep all necessary records.
- Notify local vendor marketing and/or service representatives if there is a need for immediate delivery of components to bring the computer systems to an operational level, even in a degraded mode.

If an alternate site is necessary, the following additional major tasks must be undertaken:

- Obtain governmental permissions and assistance as necessary. For example, you are likely to need building permits to construct even a temporary site.
- Coordinate moving of equipment and support personnel into the alternate site. Be sure to hire the services of a dependable security firm to protect company assets, because looting is to be expected.
- Bring the tape backups from off-site storage to the alternate site.
- As soon as the hardware is up to specifications to run the operating system, load software and run necessary tests. One of the best solutions is to record recent images of server hard drives for a fast restore.
- Prepare backup materials and return these to the off-site storage area.
- Coordinate client activities to ensure that the most critical jobs are being supported as needed.
- Be sensitive to the employees involved in the relocation. For example, in a natural disaster, employees might be suffering from the death of friends or loved ones, or they might be without a home.
- As production begins, ensure that periodic backup procedures are being followed and materials are being placed in off-site storage periodically.
- Keep administration and clients informed of the status, progress, and problems.

Chapter Summary

- Even in normal day-to-day operations, data can become corrupted or lost. Being able to restore data and applications quickly is essential.
- Factors to consider when you purchase a backup unit include the amount of data you need to back up, whether your backup software supports the unit, and the amount of money you want to spend.
- Optical disks are a good backup solution for moving infrequently used data from fast, expensive hardware (hard disks) to slower, less expensive media. However, optical disks are a poor solution if you want to back up large amounts of data quickly.
- In determining the capacity of a backup medium, you must determine the type of files the vendor used to test the drive and the type of test the vendor ran, because many capacity claims are based on using totally compressible files. These tests don't take into account that many file types can't be compressed—so the tests don't simulate real-world scenarios.
- For a large-scale enterprise, you should probably investigate an automated backup technology such as tape libraries that use autoloaders to change tapes in a largely unattended fashion. Automated tape libraries have an extensive online retention period (OLRP). At the highest level, tape libraries can support dozens of drives and hold several hundred cartridges. At that level, you will probably not configure the tape solution yourself; you tell vendor representatives what your objectives are and they will configure the library for you or provide guidance.
- There are a number of tape media standards. The most common for smaller tape devices is probably the Quarter Inch Cartridge (QIC), which as the name implies is a quarter inch in width. For server use, you probably won't see QIC and QIC Wide, which is actually .315 inches (8 mm) wide, in common use anymore.
- A Travan drive can reach compressed capacity of about 20 GB and may be useful for home, SOHO, or small departmental backups. Because Travan drives accept the QIC format, they are usually read compatible with preceding QIC standards.
- Travan NS (network series) is the most recent implementation of the Travan standard and addresses two main issues: hardware compression and fast verification.
- The mechanics of reading and writing to Travan, QIC, and similar format tapes are analogous to the way audio signals are written to a standard audio cassette. The mechanics of writing data to a DAT, in contrast, are similar to the way video signals are written to a video cassette. Rather than the tape being moved linearly across a static head, with DAT the tape is moved across an angled, rotating head.
- AIT is a Sony invention that makes backup and restore operations faster because of an optional Memory in Cassette (MIC) chip that is able to quickly locate which of the 256 tape partitions contain the data.

- Depending on the drive and media used, the DLT format allows up to 70 GB of compressed data to be stored on one rather large tape, which, unlike 8 mm or 4 mm helical scan technology, passes linearly over a fixed head. Like Travan NS technology, DLT drives can simultaneously read and write, allowing them to perform extremely well, and in some cases even better than 8 mm helical scan tape.
- The Ultrium tape format uses a single reel that stores more tape for large capacity—up to 1.6 TB per cartridge at a transfer speed up to 320 MBps. Ultrium tapes also have a memory chip that transmits its characteristics over an RF signal.
- Just like a VCR, a tape drive requires regular cleaning of the components that come into contact with the media. Without proper cleaning, tape backups can lose integrity when the heads have difficulty reading or writing through the "gunk" that forms over the heads. Some devices include automatic cleaning capability; otherwise, you will have to either procure cleaning tapes or use the manual cotton swab and cleaning solution method.
- Because tape is a magnetic medium, bulk erasing wipes out the existing data. With current tape technology, do not bulk erase. Tapes now come preformatted, and bulk erasure will remove important markings from the tape.
- Tape software offers utilities to re-tension the tape, which fast-forwards to the end without reading or writing data, and then rewinds all the way to the beginning again. This process makes tension even throughout the tape. Uneven tape tension can affect read/write reliability.
- On the whole, lower-cost tape drives, and some disk-based devices intended for use with stand-alone PCs, come bundled with backup software that provides only basic functionality. More sophisticated backup software is considerably more expensive.
- Secure backups in a locked, fireproof cabinet. You can also apply password protection to the backups and require that only the Administrator or Owner accounts can access it. You can also keep backups off-site at a service bureau.
- A normal backup copies all selected files and clears their archive bit.
- An incremental backup backs up only those files that have been created or changed since the last normal or incremental backup, which can reduce the amount of time that is required to complete the backup process. It marks files as having been backed up.
- A differential backup copies files that have been created or changed since the last normal or incremental backup, which can reduce the amount of time that is required to complete the backup process. It does not mark files as having been backed up.
- A copy backup copies all selected files, but it does not mark each file as having been backed up. Copying is useful to back up files between normal and incremental backups because it does not affect other backup operations.

- A daily backup copies all selected files that have been modified on the day that the daily backup is performed. The backed-up files are not marked as having been backed up.
- The Grandfather-Father-Son (GFS) backup strategy backs up incrementally Monday through Thursday, with a normal backup on Friday. Once a month, another normal backup is made.
- The six-cartridge backup strategy is similar to the GFS strategy, except that you do not use monthly backups and you have a two-week tape history.
- The Tower of Hanoi backup strategy requires only five tapes and creates a history that extends at least 32 days.
- Keeping accurate backup records is essential for locating backed-up data quickly, particularly if you have accumulated a large number of backup cartridges. Thorough records include cartridge labels, catalogs, online log files, and log books.
- To provide both load balancing and redundancy, a group of computers (a cluster) can cooperate in providing services. This cooperation is managed by clustering software that provides a service to clients in a client-server environment. If one node fails, the applications on the failed node are available on the other node. Throughout this process, client communications with applications usually continue with little or no interruption.
- For hardware failure recovery, having a number of spare parts available will save the time needed to order failed items from the original vendor. Some vendors, such as Intel, sell spare-parts kits comprised of the most critical system components.
- Many systems also include a hot spare that is connected to the hard disk host adapter along with the other hard disks, but is dormant until another device in the drive array fails.
- System monitoring software can detect problems such as failing hard drives, fans, power supplies, and high temperature.
- Heat is destructive to electrical components, and dust magnifies heat problems. Therefore, schedule a regular maintenance program to dust out servers.
- The combined heat of several drives in a single server can damage electrical components leading to drive failure. Be sure to provide ventilation to the drives.
- It is essential to assemble a disaster recovery team and develop a disaster recovery plan. Practical implementations will vary from one organization to another.
- In preparation for large-scale recovery, consider alternate sites. They can range from simple data collection and warehousing services designed to return data as needed to continue business operations, to full-scale facility duplication (hot site) designed to replicate all of the hardware, software, and data infrastructure necessary to assume all business functions in the event of a primary facility failure. Hot sites are very expensive to maintain, as they have nearly a complete replica of all computers in the central facility.

- Hot site recovery is appropriate for a computer operation that has more than a 24-hour outage tolerance.
- Many service bureaus augment their hot site service offering with a cold site feature. This is a facility designed to receive computer equipment. All power, water, air conditioning, raised floor, and other items requiring a long lead-time to acquire, install, and house a computer center are in place.
- Carefully consider alternate site plans, whether hot sites or cold sites. Without alternate site plans, the tendency is to make up a plan as you go, which leads to confusion and delays.

KEY TERMS

active cluster — A clustering method in which all nodes perform normal, active functions and then perform additional functions for a failed cluster member.

Advanced Intelligent Tape (AIT) — A Sony invention that uses an optional Memory in Cassette (MIC) chip on 8 mm tape that is able to quickly locate which of the 256 tape partitions contain the backed-up data.

autoloaders — Robotics inside a tape library that automatically swap tapes in and out of drives.

backup window — The optimal period of time in which you can perform a backup, usually when most files are closed.

bulk erase — Removal of data from magnetic tape using a large magnet. This is not a standard practice anymore, because most tapes require special markings that bulk erasure removes.

cold site — A disaster recovery facility designed to receive computer equipment. All power, water, air conditioning, raised floor, and other items requiring a long lead-time to acquire, install, and house a computer center are in place.

copy backup — A backup that copies all selected files, but does not mark each file as having been backed up. Copying is useful to back up files between normal and incremental backups because it does not affect other backup operations.

daily backup — A backup that copies all selected files that have been modified on the day that the daily backup is performed. The backed-up files are not marked as having been backed up.

differential backup — A backup that copies files that have been created or changed since the last normal or incremental backup, which can reduce the amount of time that is required to complete the backup process. It does not mark files as having been backed up.

digital audio tape (DAT) — Originally a high-fidelity digital recording format, now used on 4 mm and 8 mm tape backups. Uses helical scanning to record data.

digital linear tape (DLT) — A digital tape recording format that allows up to 70 GB of compressed data to be stored on one rather large tape, which, unlike 8 mm or 4 mm helical scan technology, passes linearly over a fixed head.

disaster recovery plan (DRP) — A comprehensive plan designed to recover an organization to productivity after a disaster.

failback — A clustering term referring to restoring resources to a node that has been offline when it comes back online.

Grandfather-Father-Son (GFS) — A backup strategy that uses three sets of tapes for daily, weekly, and monthly backups, retaining three months of data.

helical scanning — A tape recording method that uses a rotating tilted head to record at an angle, allowing a higher-density recording format on the tape.

Hierarchical Storage Management (HSM) — A storage management strategy in which infrequently used data is moved from expensive hardware (hard disks) to less expensive media such as optical disks or magnetic tape.

hot site — A location containing computers and necessary peripheral equipment that may be occupied or utilized by a subscriber immediately after a disaster declaration to restore its own systems, applications, and data.

hot spare — A hard disk connected to the host adapter along with the other hard disks. It is dormant until another device in the drive array fails. At that time, the hot spare is usually placed into service automatically.

incremental backup — A backup of only those files that have been created or changed since the last normal or incremental backup, which can reduce the amount of time that is required to complete the backup process. It marks files as having been backed up by setting the archive bit.

Linear Tape Open (LTO) — A collaborative technology effort headed up by HP, IBM, and Seagate to provide extremely high tape capacity and restore capability.

Linear Tape Open - Cartridge Memory (LTO-CM) — Memory in Ultrium tapes that transmits tape characteristics using radio frequency (RF) signals.

normal backup — A backup that copies all selected files and clears the archive bit for each one.

online retention period (OLRP) — References how far back in time a tape library can restore from tape without manual intervention.

optical disk — Any disk written and read by laser, including CD-R, CD-RW, DVD, and so forth.

passive cluster — A clustered server with identical services as its failover partner. A passive cluster partner remains in an idle node state until such a time as the primary node fails.

Quarter Inch Cartridge (QIC) — A common tape format that is a quarter inch wide. A variant of QIC is the QIC Wide format, which is .315 inches (8 mm) wide. QIC cartridges are generally not sufficient for server purposes.

re-tension — Fast-forwards tape to the end without reading or writing data, and then rewinds all the way to the beginning again. This process makes tension even throughout the tape.

six-cartridge backup — Similar to a GFS backup strategy, but with a two-week history.

static load balancing — A clustering technology in which a cluster member remains idle until a failure occurs.

tape library — A self-contained tape backup solution that is preloaded with several tapes. Most tape libraries include autoloaders to swap tapes.

Tower of Hanoi backup — A tape strategy that requires relatively few tapes and backs up a daily history of 32 or more days.

Travan — Created by Imation, a QIC-based standard capable of up to 20 GB compressed capacity.

Travan NS — A Travan format that can use hardware compression and fast data verification.

Ultrium — A tape format that offers a native capacity of up to 800 GB and data transfer of 80–160 MBps. Ultrium tapes use a single reel cartridge that makes room for more tape and less mechanics. Ultrium cartridges can contain memory right on the cartridge that stores a redundant file log and user-defined information.

Review Questions

1. What is the minimum number of tapes required in a GFS tape backup strategy?
   a. 12
   b. 24
   c. 48
   d. 60
2. You wish to back up all of the files on a hard disk in your computer. You have just replaced your old, single-media tape device with a new, 12-tape automated tape library system. When you attempt to initiate the backup, the process fails. What should you do?
   a. Install the driver for the new device.
   b. Restore the catalog from the old device, and then try the backup again.
   c. Make sure that you use only new blank tapes for the backup.
   d. Reboot the computer.
3. Your company's Internet servers are becoming overloaded due to an increase in commerce traffic to your web site. You decide to implement a clustering solution. What kind of clustering model should you implement in order to provide the desired load balancing?
   a. active cluster
   b. passive cluster
   c. disruptive cluster
   d. peanut cluster

4. Which of the following are necessary for disaster recovery?
   a. hot site
   b. hot-swap implementation
   c. data backup
   d. fire drills
5. Which backup media will provide the greatest storage capacity?
   a. DAT
   b. DLT
   c. Travan
   d. Ultrium
6. You want to have a separate location prepared to transfer all data management services immediately upon the failure of your network operations center. What type of site do you set up?
   a. alternate site
   b. backup storage site
   c. hot site
   d. cold site
7. Why should one or more backups be stored off-site?
   a. for security reasons
   b. because backup data is not usually needed quickly
   c. so it is available to a remote site
   d. so that primary site disasters will not affect the data
8. Which of the following are not part of a disaster recovery plan?
   a. handling legal, financial, and insurance issues
   b. performing a tape backup immediately after disaster strikes
   c. locating interim facilities to restart your business
   d. recovering PCs, LANs, and midrange systems
9. Which backup method provides the fastest backup time?
   a. Grandfather, Father, Son
   b. six cartridge
   c. Tower of Hanoi
   d. incremental backup
10. Which of the following utilizes a form of memory inside the tape? (Choose all that apply.)
    a. AIT
    b. DAT

c. Ultrium
d. QIC

11. Which of the following backs up data and clears the archive bit?
    a. normal backup
    b. differential backup
    c. copy backup
    d. daily backup
12. Which of the following Travan tape formats offers the highest compressed capacity?
    a. Travan-8
    b. Travan NS-8
    c. Travan TR-4
    d. Travan NS-20
13. Which backup strategy provides the least expensive media allocation?
    a. Grandfather, Father, Son
    b. Tower of Hanoi
    c. normal daily backup
    d. RAID-1
14. Which SNMP component should be configured in order to report errors, security breaches, or other event notification?
    a. SNMP trap
    b. SNMP Management Information Base
    c. SNMP community name
    d. SNMP host name
15. Why should spare network components be kept on-site?
    a. to configure new computers on the network quickly
    b. to provide for efficient replacement of equipment as part of a maintenance cycle
    c. to allow for efficient replacement of failed components of similar types
    d. to allow for efficient replacement of failed components of dissimilar types
16. For large-scale recovery of failed components, what site strategy is most appropriate?
    a. hot site
    b. cold site
    c. no alternate site; keep required spare components at primary site
    d. no alternate site; implement an active clustering solution

17. A _____ is a facility that has no equipment of its own, but has all the necessary facilities and environmental controls to accept server equipment.
    a. hot site
    b. cold site
    c. data center
    d. service bureau
18. For recovery from a primary facility failure with less than 24-hour downtime tolerance, what alternate site strategy is most appropriate?
    a. hot site
    b. cold site
    c. no alternate site; implement an active clustering solution
    d. no alternate site; implement a passive clustering solution
19. You want administrators to be notified of system failures as soon as possible, regardless of the time of day the failure occurs. You should implement a(n):
    a. SNMP management system
    b. tape backup strategy
    c. Management Information Base
    d. remote notification system
20. You have purchased hard disks for replacement of existing units if and when the existing units fail. You wish to be able to replace these hard disks without losing any system availability. Which elements are required to achieve this?
    a. hot-swap device capability
    b. active clustering
    c. passive clustering
    d. hot site
    e. cold site

HANDS-ON PROJECTS

Project 10-1

In this project, you will create a list of spare server components that could be found in a locked storage cabinet.

Compile a parts list for quick restoration of a server in the event of a non-system-wide failure.

Project 10-2

In this project, you will fill out a planning chart to implement a Grandfather-Father-Son (GFS) backup strategy.

You are implementing a GFS backup strategy. You desire to have one unique backup tape for each week in a month, with one unique tape for each quarter of the year. Daily and monthly tapes can be reused. Fill in the "First Quarter" chart below with the tape used. Tapes are marked with the number of the tape preceded by the day (Monday, Tuesday, etc.), month (m), or quarter (Q). Assume four weeks per month. The third week is filled in for you.

Monday 1	Tuesday 1	Wednesday 1	Thursday 1	Friday 3

Project 10-3

In this project, you will create a Tower of Hanoi backup strategy chart.

You wish to implement a Tower of Hanoi backup strategy using only four media sets. You want to maintain the greatest amount of historical backups possible. Fill in the first 14 days in the chart below with the proper allocation of backup sets for each day's backup. Label the sets A, B, C, and D.

Day	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Set														

Project 10-4

In this project, you will learn more about how a tape library works.

1. Using your Internet browser, access *www.storagetek.com*. StorageTek is a leading manufacturer of enterprise backup solutions.

2. Click the **products** link, and access information about the StorageTek L20 Tape library. This is one of StorageTek's smallest tape libraries, but it includes an informative video presentation.
3. Click a link that allows you to view the product video. If you do not have the QuickTime plug-in, follow the directions to download QuickTime, or download it from *www.quicktime.com* and then view the video.
4. Close the browser window that displayed the video, but leave the main browser open.
5. In the browser, click again on the **products** link.
6. Browse to find information on a high-end tape library such as the L700.
7. On a separate sheet of paper, write down the product name, how many drives it has, and how many tapes it can store.
8. Leave the browser open for the next project.

Project 10-5

In this project, you will learn more about disaster recovery.

1. In your web browser, access *www.davislogic.com*. Davis Logic specializes in contingency planning and disaster recovery.
2. On the left side of the Davis Logic web page, click the **Disaster Recovery** link. Notice several publications that would be excellent research guides in planning for disaster recovery.
3. On the Disaster Recovery page, scroll down the page and look for a bulleted list of possible disaster events. In this book, we have specified only a few types of disasters (such as natural disasters) that could make your site unavailable. Write down some other types of events that could constitute a disaster.
4. Close the web browser.

Project 10-6

In this project, you will perform a tape backup.

1. Verify that you have the following:
   - A server with Windows 2000
   - A tape drive
   - Tape media compatible with the tape drive
2. Log on as Administrator or some other user account that has membership in the Backup Operators group. Ask your instructor for assistance if necessary.

3. If you have not already inserted the tape into the drive, do so now. It may take a minute or more for the tape to orient itself in the drive.
4. Using Windows Explorer, create a folder at the root of the drive. Name the folder with your initials, such as "JIC."
5. Open the folder and create a Notepad text file by right-clicking in an empty space in the new folder window, clicking **New**, and then clicking **Text Document**. Name the file **MyText**.
6. Double-click the **MyText** file to open it, and type **This is the original text**. Save and exit the file.
7. Click **Start**, point to **Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
8. The Backup application starts. You can use one of the wizards at another time to see what they do. In this project, however, we will manually back up the file you created in steps 5 and 6. In the Backup application, click the **Backup** tab.
9. Expand the drive on which the folder you created is located by clicking the "+" sign next to the drive. Locate the folder and place a check in the box next to the folder.
10. At the bottom of the interface, verify that the backup destination is the appropriate tape drive. In most cases, it will be Travan.
11. In the Backup media or file name drop-down list, select **New Media**.
12. Click **Start Backup**, and a Backup Job Information dialog box opens. The backup interface should resemble Figure 10-19.

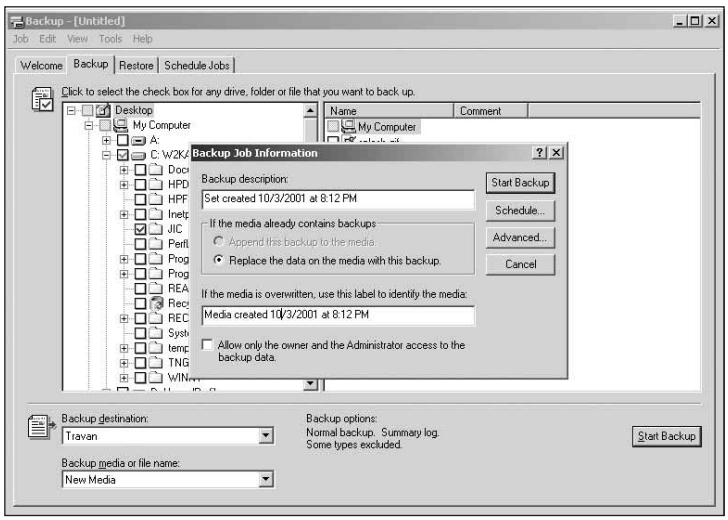

Figure 10-19 Preparing the backup job

13. Optionally, click the checkbox **Allow only the owner and the Administrator access to the backup data**.
14. Click the **Advanced** button. The Advanced Backup Options dialog box opens. Click the options **Verify data after backup** and **If possible, compress the backup data to save space**.
15. Under Backup Type, leave the setting at Normal, but click the drop-down list to look at the other backup options that are available.
16. Click **OK**. Then click the **Start Backup** button in the Backup Job Information dialog box.
17. You might see a message asking if you want to use the media detected in the tape drive. If so, click **Yes**. This does not overwrite the data; it is only a precaution to verify that you know a tape is in the drive.
18. Allow the backup to proceed. If you are prompted to overwrite the tape data, answer **Yes**. The backup may take a minute or two depending on the drive and media type.
19. When the backup is finished, do *not* click the Close button. Click the **Report** button. (Don't worry, you can look at the report later if you accidentally clicked Close.)
20. Observe the type of data that was logged.
21. Close the report, and then close the Backup utility and Windows Explorer.

Project 10-7

In this project, you will change the file you created earlier, and then restore it from tape.

1. Using Windows Explorer, access the folder and file you created in Project 10-6.
2. Open the file, and at the beginning of the file, type **This file is corrupt** as a simulation of a file you might need to restore.
3. Save and close the file.
4. Open the tape backup software (click **Start**, point to **Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**).
5. Click the **Tools** menu, and then click **Options**.
6. Click the **Restore** tab. By default, the restore process will not overwrite the file on your hard disk. In this case, the file on the hard disk is not missing; it is corrupt. Select the **Always replace the file on my computer** option, and then click **OK**.
7. Click the **Restore** tab in the main interface.

8. In the right pane, you should see the name for the media you created in Project 10-6. In the left pane, expand the "+" signs until you locate the folder that contains your file. Select that folder, and notice that the file appears in the right pane as well, similar to Figure 10-20.

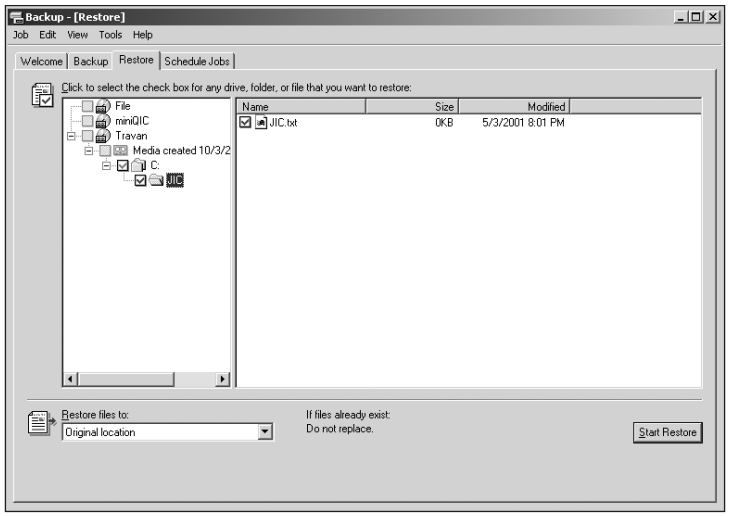

Figure 10-20 Locate the file you want to restore

9. Place a check in the box next to the folder in the left pane, and you will see that the file is automatically selected.
10. Click the **Start Restore** button.
11. A Confirm Restore dialog box appears. Click **OK** to proceed.
12. The file is restored from tape.
13. The Restore Progress dialog box should indicate that one file was processed.
14. Click the **Report** button to review what took place.
15. View the report, and when finished, close the report and then close the Backup utility.
16. Using Windows Explorer, browse to your folder and open the file. The text "This is the original text" should appear instead of "This file is corrupt." This indicates that you successfully overwrote the "corrupt" file with a known good file from backup.
17. Close Windows Explorer.

CASE PROJECTS

1. You have been asked to implement a backup strategy for your company. You currently have four administrative personnel responsible for all network support and configuration, including yourself. There is a strong desire to keep any backup processes as simple as possible. The budget for this project is less a concern than is the proper availability of data for recovery and the minimizing of the complexity of the process. Your environment consists primarily of users who store financial data on three servers. The total amount of data to be backed up is slightly over 2 GB. Describe the backup strategy that you will propose.
2. You work for a web services company, providing site hosting services to several hundred clients. What disaster recovery solutions will help you to maintain the high availability that your customers expect?